PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



# Creating an Amazon S3 bucket

**Introduction**

In this section, you will create an Amazon S3 bucket to hold the output from your worker nodes.

**Task 6.8: Create an Amazon S3 bucket**

To create an Amazon S3 bucket to hold your image files:

| Step | Action |
|---|---|
| 6.8.1 | On the **Services** menu, click **S3**. |
| 6.8.2 | Click **Create Bucket**. |
| 6.8.3 | In the **Bucket Name** box, type **image-bucket-**<*number*><br><br>**Note:** Your bucket name must be unique; so, change the number to generate a unique bucket name, such as *image-bucket-1029*. |
| 6.8.4 | To accept the default region and create your S3 bucket, click **Create**.<br><br>You are now ready to launch worker nodes within an Auto Scaling group. |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved.

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



# Launching Worker Nodes

**Introduction**

In this section, you will create an Auto Scaling group of worker nodes to process your work. After you successfully test the initial node, you will add Scaling Policies to automatically expand the size of the Auto Scaling group.

**Task 6.9: Create a launch configuration**

To create an Auto Scaling group that responds to a need to scale out by launching new instances that use your AMI:

| Step | Action |
|---|---|
| 6.9.1 | On the **Services** menu, click **EC2**. |
| 6.9.2 | In the navigation pane, click **Launch Configurations**. |
| 6.9.3 | Click **Create Auto Scaling group**. |
| 6.9.4 | Click **Create launch configuration**. |
| 6.9.5 | Click **My AMIs** tab. |
| 6.9.6 | Click **Select** to use your **Worker Image** AMI. |
| 6.9.7 | To accept the default instance type, click **Next: Configure details**. |
| 6.9.8 | On the **Create Launch Configuration** page, specify the following settings:<br>• In the **Name** box, type **Workers**<br>• For **IAM role**, select **BatchProcessing**. |
| 6.9.9 | Click **Advanced Details.** |
| 6.9.10 | Next, you will supply a user data field to the Amazon EC2 Launch Configuration so that the instances it creates are configured to automatically run the image conversion process. Because these instances will be generated from the **AMI** you created earlier, they will already have the ImageMagick and image processing software installed on them when they are launched.<br><br>In the **User data** section, verify that the **As text** option is selected.<br><br>Using the command reference file that you downloaded earlier, copy the user data text for this procedure into the **User data** box. |
| 6.9.11 | Select the option **Assign a public IP address to every instance**. |
| 6.9.12 | Click **Next: Add Storage**. |
| 6.9.13 | Click **Next: Configure Security Group** to accept the default storage setting. |
| 6.9.14 | Select the **Select an existing security group** option and then select the **BatchProcessing** security group you created earlier. |
| 6.9.15 | Click **Review**. |
| 6.9.16 | Click **Create launch configuration**. |
| 6.9.17 | Select the acknowledgement check box and then click **Create launch configuration** to accept the key pair generated by qwikLABS. |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved. 

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



## Launching Worker Nodes, Continued

**Task 6.10: Create an auto scaling group**

To create an Auto Scaling group:

| Step | Action |
|---|---|
| 6.10.1 | On the **Configure Auto Scaling Group** details page, specify the following settings:<br>• In the **Group name** box, type **worker-group**<br>• In the **Network** drop-down list, click **Lab VPC**.<br>• Click inside the **Subnet** box, and select either **Public Subnet 1** or **Public Subnet 2**. |
| 6.10.2 | Click **Next: Configure scaling policies**. |
| 6.10.3 | Select the **Use scaling policies to adjust the capacity of this group** option and change the policy to "Scale between **1** and **4** instances." |
| 6.10.4 | In the **Increase Group Size** section, change **Take the action** to Add **1** instances. **And then wait** to **180** seconds before allowing another scaling activity. |
| 6.10.5 | Click **Review**. |
| 6.10.6 | Click **Create Auto Scaling group** and then click **Close**.<br><br>Your Auto Scaling group has been configured to run only a single instance now, and no alarms have been attached to its scaling policies. After you verify that the worker node is functioning correctly, you can create alarms to automatically adjust the number of worker nodes. |
| 6.10.7 | In the navigation pane, click **Instances**.<br><br>A new instance will start now. Wait until the Instance State changes to *running* and the Status Checks shows *2/2 checks passed.* |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved. 

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



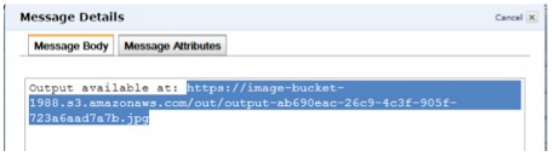
# Dispatching Work and Viewing Results

**Introduction**

In this section, you will use the **SQS Dashboard** to view the output of your last message, and use a message producer script to add more messages to your Amazon SQS input queue in order to use CloudWatch SQS metrics in the next section.

**Task 6.11: Check your queue for messages**

To verify that the message you sent to the input queue earlier has been processed and moved to the output queue:

| Step | Action |
|---|---|
| 6.11.1 | On the **Services** menu, click **SQS**.<br><br>Confirm that there is **1 Message Available** in your **output** queue. If the message is still in your input queue, click **Refresh** periodically. It may take a few minutes after your worker instances were created for the message to move to the **output** queue from the **input** queue.<br><br>**Note:** If there is no message in your **output** queue:<br>• Ensure that your queues are named **input** and **output** (in lowercase).<br>• Ensure that your **BatchProcessing** role has granted full permissions for Amazon SQS and Amazon S3.<br>• Ensure that your worker node is using the **BatchProcessing** IAM role (defined in the Launch Configuration).<br>• Ensure that your worker node is running.<br>Ask your instructor for assistance in successfully running your worker node. |
| 6.11.2 | Select the **output** queue. Ensure that the **input** queue is not selected. |
| 6.11.3 | In the **Queue Actions** drop-down list, click **View/Delete Messages**. |
| 6.11.4 | Click **Start Polling for Messages**. |
| 6.11.5 | Find your message and click **More Details** to view the message body.<br><br>The message contains the output. To view the image, select and copy the URL portion of the message. Then paste the link into a new browser window.<br><br> |
| 6.11.6 | Click **Close** to close the **Message Details** window. |
| 6.11.7 | Click **Close** to close the **View/Delete Messages in output** window. |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved.

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



# Dispatching Work and Viewing Results, Continued

**Task 6.12: Connect to your worker instance and run the message production script (Windows)**

**Note:** This section is for **Windows** users only. **Mac** or **Linux** users, skip to Task 6.13.

To run the message production script using Windows:

| Step | Action |
|---|---|
| 6.12.1 | Launch **PuTTY**. |
| 6.12.2 | On the **Services** menu, click **EC2**. |
| 6.12.3 | In the Navigation pane, click **Instances**. |
| 6.12.4 | Select your worker instance.<br><br>**Note:** Your worker instance will be unnamed; however, it will be the only other instance available other than your **Master** instance. |
| 6.12.5 | Copy the **Public DNS** name of your worker instance from the EC2 Dashboard and paste it into the **Host Name** box in your **PuTTY** client. |
| 6.12.6 | In the **Category** list, expand **SSH**, and then click **Auth**. |
| 6.12.7 | In the **Private key file for authentication** box, browse to the **PPK** file that you downloaded in **Task 6.3**. |
| 6.12.8 | Click **Open** to initiate the connection, and log in as **`ec2-user`**. |
| 6.12.9 | Run the following command:<br><br>**`sudo python MessageProducer.py`**<br><br>**Note:** You should receive a message that confirms that the message was sent to your input queue. If you do not receive this confirmation, notify your instructor. If attempts to run the script continue to fail, you can emulate this script's behavior by using steps 6.7.6–6.7.9 to send at least 35 messages from your input queue.<br><br>The purpose of this script is to generate a long backup of messages in your input queue, so that the Amazon CloudWatch metric associated with visible queue messages will be available in the next task. |
| 6.12.10 | Proceed to **Task 6.14: Monitoring the Cluster**. |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved.

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



## Dispatching Work and Viewing Results, Continued

**Task 6.13: Connect to your worker instance and run the message production script (Mac and Linux)**

**Note:** This section is for **Mac** or **Linux** users only.

To run the message production script using a Mac or Linux machine:

| Step | Action |
|---|---|
| 6.13.1 | On the **Services** menu, click **EC2**. |
| 6.13.2 | In the Navigation pane, click **Instances**. |
| 6.13.3 | Select your worker instance.<br><br>**Note:** Your worker instance will be unnamed; however, it will be the only other instance available other than your **Master** instance. |
| 6.13.4 | Copy the **Public DNS** value of your worker instance from the EC2 Dashboard. |
| 6.13.5 | Run the following command to connect to the Amazon EC2 instance:<br><br>`ssh -i <location-of-pem> ec2-user@<host-IP>`<br><br>Substitute the Public DNS value for *<host-IP>*. |
| 6.13.6 | Run the following command:<br><br>`sudo python MessageProducer.py`<br><br>**Note:** You should receive a message that confirms that the message was sent to your input queue. If you do not receive this confirmation, notify your instructor. If attempts to run the script continue to fail, you can emulate this script's behavior by using steps 6.7.6–6.7.9 to send at least 35 messages from your input queue.<br><br>The purpose of this script is to generate a long backup of messages in your input queue, so that the Amazon CloudWatch metric associated with visible queue messages will be available in the next task. |
| 6.13.7 | Proceed to **Task 6.14: Monitoring the Cluster**. |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved. 

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



# Monitoring the Cluster

**Introduction**

You can now use CloudWatch to monitor your cluster. You will define a CloudWatch alarm for use with Auto Scaling policies.

**Task 6.14: Create an alarm**

To create an alarm that triggers your Auto Scaling group:

| Step | Action |
|---|---|
| 6.14.1 | On the **Services** menu, click **CloudWatch**. |
| 6.14.2 | Click **Browse Metrics.** |
| 6.14.3 | To expand the drop-down list, click **Browse Metrics** again. |
| 6.14.4 | In the **Browse Metrics** drop-down list, click **SQS**.<br><br>**Note:** If the **SQS Metrics** header is not visible, return to your **input** queue and ensure that there are messages queued for processing. This will trigger metrics to be sent to CloudWatch after a few minutes. The **Visible** count may take a bit longer to appear. If the list of messages waiting in your input queue is lower than 25, run the message producer script again. |
| 6.14.5 | Select the **ApproximateNumberOfMessagesVisible** metric for the **input** queue. |
| 6.14.6 | On the **Tools** menu in the lower-right corner, click **Create Alarm**. |
| 6.14.7 | On the **Alarm Threshold** page, specify the following settings:<br>• In the **Name** box, type **long-queue**<br>• In the **Description** box, type **Queue too long**<br>• Enter **10** so that the threshold is:<br>**Whenever:** ApproximateNumberOfMessagesVisible<br>**is:** >= 10<br>**for:** 1 consecutive period(s) |
| 6.14.8 | In **Actions**, click **Delete** on the existing Notification Action. |
| 6.14.9 | Click **+ AutoScaling Action**, and specify the following settings:<br>• For **From the group**, select **worker-group**.<br>• For **Take this action**, select **Increase Group Size – Add 1 instance**. |
| 6.14.10 | Click **Create Alarm**.<br><br>**Note:** These configuration settings will automatically **add 1 instance** to your Auto Scaling group called **worker-group** whenever the **input** queue has **10 or more messages visible** within a **1 minute period**. |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved.

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



## Monitoring the Cluster, Continued

**Task 6.15: Test your Auto Scaling group**

To test and verify that your Auto Scaling group adds and removes instances:

| Step | Action |
|---|---|
| 6.15.1 | Switch to your **CloudWatch Dashboard**. Your **long-queue** alarm should appear within the **Alarm Summary** section. |
| 6.15.2 | Refresh the **Alarm Summary** section until it displays **"You have 1 alarm in ALARM state…"** |
| 6.15.3 | Switch back to the **EC2 Dashboard**, and select **Instances** in the navigation pane.<br><br>Within a few minutes, you should be able to see the additional instances that have been launched by your Auto Scaling group. As a result, the messages in your queue will be processed more quickly. |
| 6.15.4 | In the navigation pane, click **Auto Scaling Groups** and select your **worker-group**.<br><br>On the **Scaling History** tab, you can verify that Auto Scaling has launched new instances. |

**Ending your lab**

When you are done, follow the instructions in Appendix B to end your lab.

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved.

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



## Optional

**Introduction** This part of the lab is a challenge wherein step-by-step instructions are not given.

**Challenge** Currently, your Auto Scaling group only scales out. Any instances that the group launches in order to handle a surplus of messages will continue to run even after those messages have been processed.

Modify your environment's settings so that your Auto Scaling group scales in once the number of messages in your input queue drops below 10. You can use the message producer script to load up your input queue again, and verify that your environment's new settings successfully scale in your instances.

**Ending your lab** When you are done, follow the instructions in **Appendix B** to end your lab.

## Conclusion

**Achievements** Having completed this lab, you are able to:

- Bootstrap an EC2 instance using User Data.
- Create an AMI from a running instance.
- Use the AWS Management Console to create an SQS queue.
- Create an Auto Scaling group with Scaling Policies based on an SQS queue.
- Use the AWS Management Console to pass messages to, and read messages from, an SQS queue.

Your results are verified by your ability to complete the lab in the allotted time. When you have completed the lab, your instructor will facilitate a discussion of your results.

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved. 

PRINTED BY: gleung@datapipe.com. Printing is for personal, private use only. No part of this book may be reproduced or transmitted without publisher's prior permission. Violators will be prosecuted.



# Appendix A: Downloading your key pair or remote desktop shortcut

**Introduction**

Access to instances used in your labs requires a secure connection using either an SSH client (for Linux instances) or Remote Desktop Connection (for Windows instances). The qwikLABS page for your lab provides access to the key pair or remote desktop shortcut used to access these instances.

**Downloading your key or shortcut**

To download your key pair or remote desktop shortcut:

| Step | Action |
|---|---|
| 1 | From the qwikLABS page in your browser, in the **Connection Details** section, click the **Download PEM/PPK** drop-down and choose either the PPK file (for PuTTY clients on Windows) or the PEM file (for Macintosh clients). |
| 2 | Save the file to your `\Downloads` folder or an alternate location. |
| 3 | If your lab requires access to a Windows instance, in the **RDP** section, click the **Download RDP File** option. |
| 4 | Save the file to your `\Downloads` folder or an alternate location. |
| 5 | In the **RDP** section of the page, make a note of the RDP **Password**. You will need it to connect to your Windows instance using the RDP shortcut. |

© 2015 Amazon Web Services, Inc. and its affiliates. All rights reserved. 